arXiv:1207.2944v1 [cond-mat.mes-hall] 12 Jul 2012

Feedback control of a solid-state qubit using high-fidelity projective measurement

D. Ristè,[1] C. C. Bultink,[1] K. W. Lehnert,[2] and L. DiCarlo[1]

[1] *Kavli Institute of Nanoscience, Delft University of Technology, P.O. Box 5046, 2600 GA Delft, The Netherlands*

[2] *JILA, National Institute of Standards and Technology and the University of Colorado and Department of Physics, University of Colorado, Boulder, Colorado 80309, USA*

(Dated: July 13, 2012)

We demonstrate feedback control of a superconducting transmon qubit using discrete, projective measurement and conditional coherent driving. Feedback realizes a fast and deterministic qubit reset to a target state with 2.4% error averaged over input superposition states, and cooling of the transmon from 16% spurious excitation to 3%. This closed-loop qubit control is necessary for measurement-based protocols such as quantum error correction and teleportation.

PACS numbers: 03.67.Lx, 42.50.Dv, 42.50.Pq, 85.25.-j

Many protocols in quantum information processing (QIP) require closing a feedback loop where coherent control of qubits is conditioned on projective measurements in real time [1]. Important examples include quantum error correction and teleportation [2], so far achieved in ion-trap [3, 4] and photonic systems [5, 6]. During the last decade, the steady development of qubit readout and universal gates needed in a quantum processor [7] has made superconducting circuits [8] a leading solid-state QIP platform. However, the simple quantum algorithms [9] and teleportation-like protocol [10] so far demonstrated fall in the category of open-loop control. Measurement is performed as the final step, following a programmed sequence of applied gates. A comparable realization of closed-loop control has been precluded by stringent requirements on high measurement fidelity and short loop delay (latency). Until recently, the available qubit coherence times bottlenecked both achievable fidelity and required speed.

For feedback control of superconducting qubits, the development of circuit quantum electrodynamics [11, 12] with 3D cavities (3D cQED) [13] constitutes a watershed. The new order of magnitude in qubit coherence times (> 10 μs) allows boosting projective-readout fidelity up to 98% [14, 15] and realizing feedback with off-the-shelf electronics. Very recently, feedback based on continuous weak measurement has sustained Rabi oscillations of a transmon qubit indefinitely [16]. Previously, this type of feedback had only been used to generate and stabilize quantum states of photons [17], ions [18], and atoms [19].

In this Letter, we demonstrate feedback control of a superconducting transmon qubit based on discrete, projective measurement. This dual type of feedback is the kind necessary for measurement-based QIP. As a first application, we demonstrate a feedback-based reset that is deterministic and fast compared to passive initialization. This feedback cools the transmon from a spurious steady-state excitation of 16% to 3% and resets qubit states with 2.4% error averaged over the Bloch sphere. These absolute errors are dominated by latency, in quantitative agreement with a model including transmon equilibration and readout errors.

The experiment employs a transmon qubit inside an aluminum 3D cavity [13]. The qubit ($\omega_{01}/2\pi = 4.889$ GHz transition frequency) couples to the cavity fundamental mode ($\omega_r/2\pi = 6.546$ GHz, coupling-limited linewidth $\kappa/2\pi = 550$ kHz) with strength $g/2\pi = 68$ MHz. The high-fidelity, projective qubit readout forming the input to the feedback loop uses homodyne detection of the qubit-state dependent cavity transmission (dispersive shift $2\chi/2\pi = -1.9$ MHz [12]). A 400 ns measurement pulse at $\omega_{RF} = \omega_r - \chi$ is applied to the cavity and the transmitted signal is then amplified by a

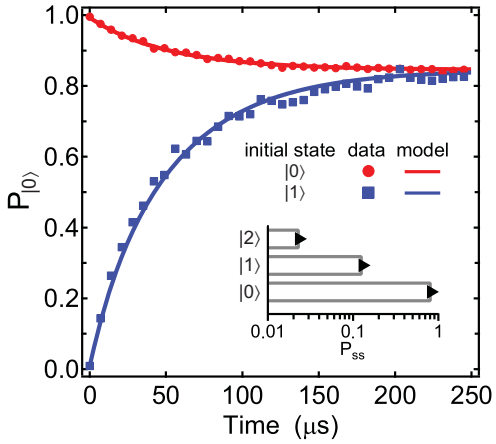

FIG. 1. (color online). Transmon equilibration to steady state. Time evolution of the ground-state population $P_{|0\rangle}$ starting from states $|0\rangle$ and $|1\rangle$. Solid curves are the best fit (including data in Fig. S4) to Eq. 1, giving the inverse transition rates $\Gamma_{01}^{-1} = 50 \pm 2$ μs, $\Gamma_{12}^{-1} = 20 \pm 2$ μs, $\Gamma_{10}^{-1} = 324 \pm 32$ μs, $\Gamma_{21}^{-1} = 111 \pm 25$ μs. From the steady-state solution, we extract residual excitations $P_{|1\rangle,ss} = 13.1 \pm 0.8\%$, $P_{|2\rangle,ss} = 2.4 \pm 0.4\%$. Inset: steady-state population distribution (bars). Markers correspond to a Boltzmann distribution with best-fit temperature 127 mK, significantly higher than the dilution refrigerator base temperature (15 mK).

Josephson parametric amplifier [20, 21] to enhance sensitivity [22]. The feedback controller is an ADwin Gold processor that samples the transmitted homodyne signal, performs 1-bit digitization to interpret the projected qubit state, and conditionally triggers a π pulse resonant with the transmon $0 \leftrightarrow 1$ transition. The 2.64 μs delay between start of the measurement and end of the π pulse, set by processing time in the ADwin, is short compared to the qubit relaxation time T_1 (see below) [23].

Our first application of feedback is qubit initialization, also known as reset [7]. To be useful in QIP, reset must be deterministic (as opposed to heralded or postselected [14, 15]) and fast compared to qubit coherence times. Obviously, the passive method, i.e., waiting several T_1 times, does not meet the speed requirement. Moreover, it can suffer from residual steady-state qubit excitations [14–16, 24], whose source remains an active area of research. The drawbacks of passive initialization are evident for our qubit, whose ground-state population $P_{|0\rangle}$ evolves from states $|0\rangle$ and $|1\rangle$ as shown in Fig. 1. The transmon dynamics are captured by a master equation model for a three-level system [22, 25]:

$$\begin{pmatrix} \dot{P}_{|0\rangle} \\ \dot{P}_{|1\rangle} \\ \dot{P}_{|2\rangle} \end{pmatrix} = \begin{pmatrix} -\Gamma_{10} & \Gamma_{01} & 0 \\ \Gamma_{10} & -\Gamma_{01}-\Gamma_{21} & \Gamma_{12} \\ 0 & \Gamma_{21} & -\Gamma_{12} \end{pmatrix} \begin{pmatrix} P_{|0\rangle} \\ P_{|1\rangle} \\ P_{|2\rangle} \end{pmatrix}. \tag{1}$$

The best fit to the data gives the qubit relaxation time $T_1 = 1/\Gamma_{01} = 50 \pm 2$ μs and 15.5% residual total excitation. Previous approaches to accelerate qubit equilibration include coupling to dissipative resonators [26] or two-level systems [27]. However, these are also susceptible to spurious excitation, potentially inhibiting complete qubit relaxation.

Feedback-based reset circumvents the equilibration problem by not relying on coupling to a dissipative medium. Rather, it works by projecting the qubit with a measurement (M_1, performed by the controller) and conditionally applying a π pulse to drive the qubit to a targeted basis state (Fig. 2). A final measurement (M_2) determines the qubit state immediately afterwards, by averaging the second half (200 ns) of a transmitted measurement pulse sampled by a data acquisition card [22]. In both measurements, the result is digitized into levels H or L, associated with $|0\rangle$ and $|1\rangle$, respectively. The digitization threshold voltage $V_{\rm th}$ maximizes the readout fidelity at 99%. The π pulse is conditioned on $M_1 = L$ to target $|0\rangle$ (scheme Fb_0) or on $M_1 = H$ to target $|1\rangle$ (Fb_1). To benchmark the reset protocol, we quantify its action on the purest states we can prepare. This step is accomplished with a preliminary measurement M_0 (initializing the qubit in $|0\rangle$ by postselection [14, 15]), followed by a calibrated pulse resonant on the transmon $0 \leftrightarrow 1$ transition to prepare $|\theta\rangle = \cos(\theta/2)\,|0\rangle + \sin(\theta/2)\,|1\rangle$. We first consider the action of reset on the basis states ($\theta = 0, \pi$). The overlap of the M_2 histograms with the targeted region (H for Fb_0 and L for Fb_1) averages at 96%, indicating the success of reset. Imperfections are more evident for $\theta = \pi$ and mainly due to equilibration of the transmon during the feedback loop. We address this point quantitatively below.

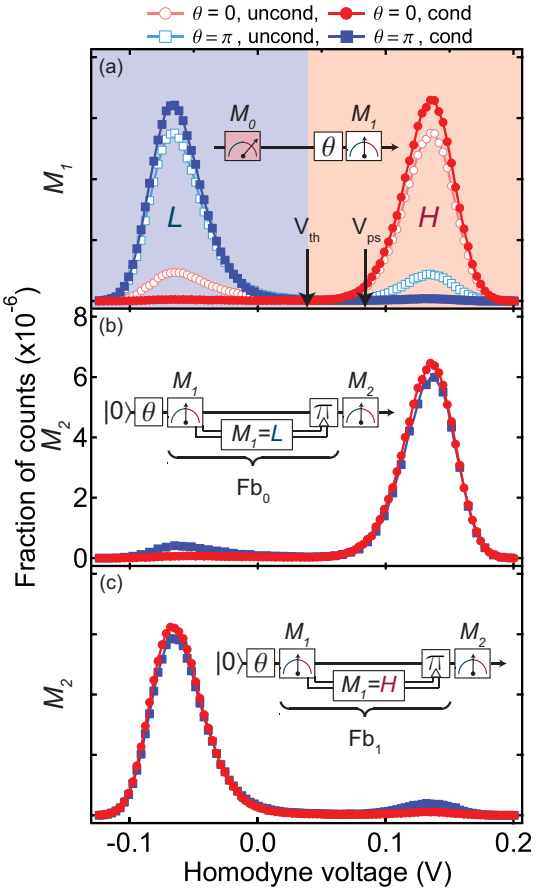

FIG. 2. (color online). Reset by measurement and feedback. (a) Histograms of 300 000 shots of M_1, with (squares) and without (circles) inverting the qubit population with a π pulse. H and L denote the two possible outcomes of M_1, digitized with the threshold $V_{\rm th}$, maximizing the contrast. Full (empty) dots indicate (no) postselection on $M_0 > V_{\rm ps}$, chosen to prepare $|0\rangle$ with 0.7% error while keeping 81% of the data. This technique purifies the steady-state mixture (see Fig. 1), preparing the input states for the feedback sequences in (b) and (c). (b) Histograms of M_2 after applying the feedback protocol Fb_0, which triggers a π pulse when $M_1 = L$. Using this feedback, $\sim$ 99% (92%) of counts fall in the H area for $\theta = 0$ (π), respectively. (c) Feedback with opposite logic Fb_1 preparing the excited state. In this case, $\sim$ 98% (94%) of counts fall in the L area for $\theta = 0$ (π).

An ideal reset function always prepares the same pure qubit state, regardless of its input. To fully quantify the performance of our reset scheme, we measure its effect on different superposition states $|\theta\rangle$. After applying Fb_0 (Fb_1), we extract $P_{|0\rangle}$ by averaging over 300 000 repetitions of the experiment (Fig. 3). Without feedback, $P_{|0\rangle}$ is trivially a sinusoidal function of θ, with contrast

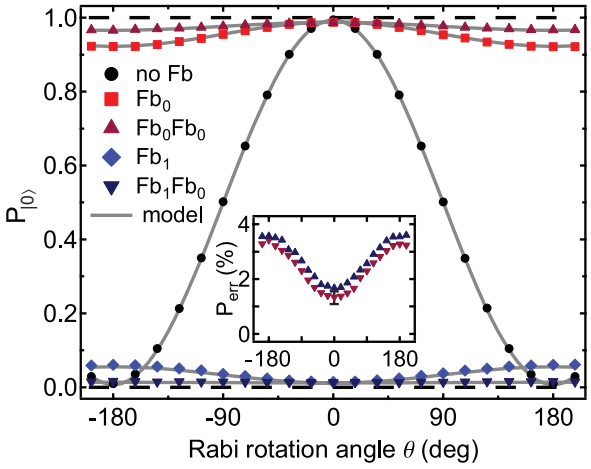

FIG. 3. (color online). Deterministic reset from any qubit state. Ground-state population $P_{|0\rangle}$ as a function of the initial state $|\theta\rangle = \cos(\theta/2)\,|0\rangle + \sin(\theta/2)\,|1\rangle$, prepared by coherent rotation after initialization in $|0\rangle$, as in Fig. 2. The cases shown are: no feedback (circles), one instance of Fb_0 (squares) or Fb_1 (diamonds), and two instances of Fb_0 (upward triangles) or Fb_0 followed by Fb_1 (downward triangles). Inset: error probabilities for two rounds of feedback, defined as $1 - P_{|t\rangle}$, where $|t\rangle \in \{0, 1\}$ is the target state. The vertical axis is calibrated with the reference states (Fig. 2) and corrected for the partial excitation between M_0 and M_1. The systematic $\sim 0.3\%$ difference between the two cases is attributed to error in the π pulse preceding the measurement of $P_{|1\rangle}$ following Fb_1 [22]. Curves: model including readout errors and equilibration [22].

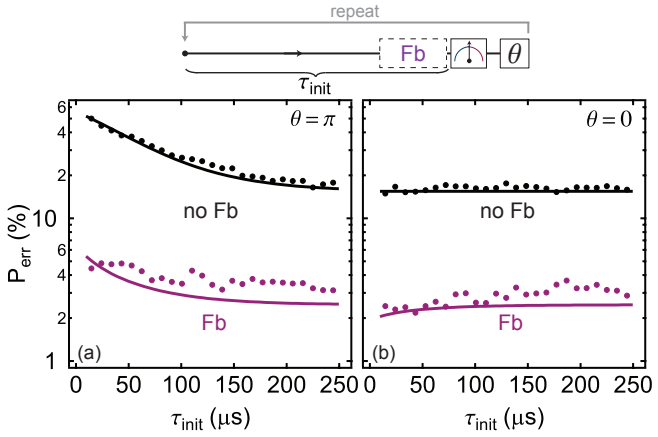

FIG. 4. (color online). Fast qubit reset. Initialization errors under looped execution of a simple experiment consisting of (a) a π pulse and (b) no pulses, leaving the qubit ideally in $|1\rangle$ or $|0\rangle$, respectively. Subsequent runs are separated by an initialization time τ_{init}. Black: initialization by waiting (no feedback). Purple: initialization by feedback, with three rounds of Fb_0 and a π pulse on the $1 \leftrightarrow 2$ transition [22]. Curves correspond to a master equation simulation assuming perfect pulses.

slightly reduced from unity due to a preparation infidelity of 0.7%. With feedback, we observe a suppression of the Rabi oscillation, with $P_{|0\rangle}$ approaching the ideal value 1 (0) for Fb_0 (Fb_1) for any input state. However, a dependence on θ remains, with $P_{\mathrm{err}} = 1 - P_{|0\rangle}$ for Fb_0 ($1 - P_{|1\rangle}$ for Fb_1) ranging from 1.2% (1.4%) for $\theta = 0$ to 7.8% (8.4%) for $\theta = \pi$. The remaining errors have two sources: mismatch between measurement result M and post-measurement state $|i\rangle$, occurring with probability p^M_{ij}, for initial state $|j\rangle$; and equilibration during the $\tau = 2.4$ μs lapse between the end of M_1 and the start of M_2, set by processing time in the controller. Transitions to $|2\rangle$ during M_1 (with probability $p_{21} = p^H_{21} + p^L_{21}$), or during τ ($\Gamma_{21}\tau$), cause leakage out of the qubit subspace, where the feedback has no action. For perfect pulses, the overall errors (equal for Fb_0 and Fb_1) are to first order:

$$P^{\theta=0}_{\mathrm{err}} = p^L_{00} + p^H_{10} + \Gamma_{10}\tau,$$
$$P^{\theta=\pi}_{\mathrm{err}} = p^H_{11} + p^L_{01} + p_{21} + (\Gamma_{01} + \Gamma_{21})\tau,$$

and weighted combinations thereof for other θ. Using the best-fit rates Γ_{ij} (Fig. 1), errors due to equilibration sum to 0.7% (6.9%) for $\theta = 0$ (π), while readout errors account for the remaining 0.4% (1.4%).

To improve reset fidelities, we concatenate two feedback cycles. The dominant error for $\theta = 0$ remains unchanged, but for $\theta = \pi$ it decreases to $P^{\theta=0}_{\mathrm{err}} + p_{21} + \Gamma_{21}\tau$, in agreement with the measured values of 1.3% and 3.4%, respectively. The consecutive application of Fb_0 and Fb_1 prepares $|1\rangle$ with similar fidelities.

The key advantage of reset by feedback is the ability to ready a qubit for further computation fast compared to coherence times available in 3D cQED [13, 28]. This will be important for ancilla qubits used for multiple rounds of error correction [29]. To test the gain in speed over passive initialization, we run a sequence multiple times, separated by a variable initialization time τ_{init}, during which we perform reset or not (Fig. 4). A measurement pulse follows the initialization period to quantify the initialization error P_{err}. The sequence itself emulates an algorithm leaving the qubit in $|1\rangle$ (or $|0\rangle$), by simply applying (or not applying) a π pulse. Without feedback, when the sequence is a π pulse [Fig. 4(a)] the qubit remains partially excited and P_{err} approaches 50% as $\tau_{\mathrm{init}} \to 0$. With no π pulse [Fig. 4(b)], P_{err} simply equals the total steady-state excitation. Feedback during τ_{init} suppresses the reset error from $\sim 32\%$ (averaged between the two cases) to 3.5% for the shortest τ_{init} (15 μs). We achieve this improvement by combining three rounds of Fb_0 with a π pulse on the $1 \leftrightarrow 2$ transition before the final Fb_0, which partially counters leakage to the second excited state [22]. In the opposite limit $\tau_{\mathrm{init}} \gg T_1$, residual excitations are reduced from 16% to 3% [30], cooling the transmon.

In conclusion, we have demonstrated feedback control of a transmon qubit using high-fidelity projective measurement and conditional operation. We have applied

this feedback to deterministically reset the qubit, starting from any superposition, with an average error of 2.4%. We have also used feedback-based initialization to loop a test sequence with more than 15-fold increase in repetition rate relative to the passive method. Reset fidelity is currently limited by transitions to higher energy levels, caused by the 127 mK transmon temperature. We estimate that by reducing this temperature below 50 mK, the suppression of Γ_{10} and Γ_{21} will decrease the reset error past the $\sim 1\%$ fault-tolerant threshold in surface coding [31]. This reduction may be achieved by a combination of infrared radiation shielding [24], use of a copper cavity [32], and improved qubit thermal anchoring. Moreover, decreasing the latency in the feedback loop, for example by using field-programmable gate arrays, will further reduce errors due to equilibration between measurement and action. While this demonstration employs feedback on a single qubit, the scheme can be extended to conditionally drive another qubit within its coherence time, realizing the feedforward [1] needed for teleportation and measurement-based error correction. Future experiments will also target the generation of entanglement by combining the demonstrated feedback scheme with cavity-based parity measurements [33, 34].

We thank V. Ranjan, J. G. van Leeuwen, and H.-S. Ku for experimental assistance, and M. Tiggelman, R. N. Schouten, and A. Wallraff for discussions. We acknowledge funding from the Dutch Organization for Fundamental Research on Matter (FOM), the Netherlands Organization for Scientific Research (NWO, VIDI scheme), the EU FP7 project SOLID, and the DARPA QuEST program.

Note added.—A parallel manuscript from École Normale Supérieure, Paris [35] reports feedback control of a transmon qubit using a Josephson parametric converter for high-fidelity projective readout.

[1] H. M. Wiseman and G. J. Milburn, *Quantum measurement and control* (Cambridge University Press, Cambridge, 2009).
[2] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000).
[3] J. Chiaverini, *et al.*, Nature **432**, 602 (2004).
[4] M. Riebe, *et al.*, Nature **429**, 734 (2004).
M. D. Barrett, *et al.*, Nature **429**, 737 (2004).
[5] T. B. Pittman, B. C. Jacobs, and J. D. Franson, Phys. Rev. A **71**, 052332 (2005).
[6] A. Furusawa, J. L. Sørensen, S. L. Braunstein, C. A. Fuchs, H. J. Kimble, and E. S. Polzik, Science **282**, 706 (1998).
[7] D. P. DiVincenzo, Fortschritte der Physik **48**, 771 (2000).
[8] J. Clarke and F. K. Wilhelm, Nature **453**, 1031 (2008).
[9] L. DiCarlo, *et al.*, Nature **460**, 240 (2009).
T. Yamamoto, *et al.*, Phys. Rev. B **82**, 184515 (2010).
A. Dewes, R. Lauro, F. R. Ong, V. Schmitt, P. Milman, P. Bertet, D. Vion, and D. Esteve, Phys. Rev. B **85**, 140503 (2012).
E. Lucero, *et al.*, arXiv:1202.5707.
[10] M. Baur, A. Fedorov, L. Steffen, S. Filipp, M. P. da Silva, and A. Wallraff, Phys. Rev. Lett. **108**, 040502 (2012).
[11] A. Blais, R.-S. Huang, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf, Phys. Rev. A **69**, 062320 (2004).
[12] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. Kumar, S. M. Girvin, and R. J. Schoelkopf, Nature **431**, 162 (2004).
[13] H. Paik, *et al.*, Phys. Rev. Lett. **107**, 240501 (2011).
[14] D. Ristè, J. G. van Leeuwen, H.-S. Ku, K. W. Lehnert, and L. DiCarlo, arXiv:1204.2479, Phys. Rev. Lett., in press (2012).
[15] J. E. Johnson, C. Macklin, D. H. Slichter, S. J. Weber, E. B. Weingarten, J. Clarke, and I. Siddiqi, arXiv:1202.5541, Phys. Rev. Lett., in press (2012).
[16] R. Vijay, C. Macklin, D. H. Slichter, K. W. Murch, R. Naik, N. Koroktov, and I. Siddiqi, arXiv:1205.5591.
[17] G. G. Gillett, *et al.*, Phys. Rev. Lett. **104**, 080503 (2010).
C. Sayrin, *et al.*, Nature **477**, 73 (2011).
[18] P. Bushev, *et al.*, Phys. Rev. Lett. **96**, 043003 (2006).
[19] M. Koch, C. Sames, A. Kubanek, M. Apel, M. Balbach, A. Ourjoumtsev, P. W. H. Pinkse, and G. Rempe, Phys. Rev. Lett. **105**, 173003 (2010).
S. Brakhane, W. Alt, T. Kampschulte, M. Martinez-Dorantes, R. Reimann, S. Yoon, A. Widera, and D. Meschede, arXiv:1206.3184.
[20] M. A. Castellanos-Beltran, K. D. Irwin, G. C. Hilton, L. R. Vale, and K. W. Lehnert, Nature Phys. **4**, 929 (2008).
[21] R. Vijay, M. H. Devoret, and I. Siddiqi, Rev. Sci. Instrum. **80**, 111101 (2009).
[22] See supplemental material.
[23] The Ramsey dephasing time $T_2^* = 2$ μs is most likely due to photon shot noise [36] and sensitivity to charge noise due to the low ratio $E_J/E_C \sim 27$ [37]. Dephasing is however not relevant for the reset application here demonstrated.
[24] A. D. Córcoles, J. M. Chow, J. M. Gambetta, C. Rigetti, J. R. Rozen, G. A. Keefe, M. Beth Rothwell, M. B. Ketchen, and M. Steffen, Appl. Phys. Lett. **99**, 181906 (2011).
[25] R. Bianchetti, S. Filipp, M. Baur, J. M. Fink, M. Göppl, P. J. Leek, L. Steffen, A. Blais, and A. Wallraff, Phys. Rev. A **80**, 043840 (2009).
[26] M. D. Reed, B. R. Johnson, A. A. Houck, L. DiCarlo, J. M. Chow, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf, Appl. Phys. Lett. **96**, 4 (2010).
[27] M. Mariantoni, *et al.*, Science **334**, 61 (2011).
[28] C. Rigetti, *et al.*, arXiv:1202.5533.
[29] P. Schindler, J. T. Barreiro, T. Monz, V. Nebendahl, D. Nigg, M. Chwalla, M. Hennrich, and R. Blatt, Science **332**, 1059 (2011).
[30] We note that $P_{|1\rangle} \approx P_{|2\rangle} = 1.6\%$ is a non-thermal distribution.
[31] D. S. Wang, A. G. Fowler, and L. C. L. Hollenberg, Phys. Rev. A **83**, 020302 (2011).
[32] We have recently measured a residual excitation of $\sim 2\%$ for a qubit at the same frequency inside a Cu cavity of the same geometry, consistent with Ref. 28.

[33] K. Lalumière, J. M. Gambetta, and A. Blais, Phys. Rev. A **81**, 040301 (2010).
[34] L. Tornberg and G. Johansson, Phys. Rev. A **82**, 012329 (2010).
[35] P. Campagne-Ibarcq, *et al.*, in preparation (2012).
[36] A. P. Sears, *et al.*, arXiv:1206.1265.
[37] J. A. Schreier, *et al.*, Phys. Rev. B **77**, 180502 (2008).

Supplement to "Feedback control of a solid-state qubit using high-fidelity projective measurement"

D. Ristè,[1] C. C. Bultink,[1] K. W. Lehnert,[2] and L. DiCarlo[1]

[1] *Kavli Institute of Nanoscience, Delft University of Technology, P.O. Box 5046, 2600 GA Delft, The Netherlands*

[2] *JILA, National Institute of Standards and Technology and the University of Colorado and Department of Physics, University of Colorado, Boulder, Colorado 80309, USA*

(Dated: July 13, 2012)

arXiv:1207.2944v1 [cond-mat.mes-hall] 12 Jul 2012

READOUT AND FEEDBACK CONFIGURATIONS

The readout frequency is set to $\omega_{\rm RF} = \omega_{\rm r} - \chi$ to maximize the phase difference in cavity transmission between qubit in $|0\rangle$ and $|1\rangle$. The JPA is tuned to $\omega_{\rm RF}$ by the combination of external flux biasing and a continuous-wave pump tone that bends the JPA resonance, as described in Ref. 1. In contrast to this previous work, here we use a single generator for pump, measurement and local oscillator in order to maximize phase stability. A phase shifter on the RF arm is used to maximize the phase-sensitive amplification [Fig. S3(a)]. Another phase shifter at the output side maximizes the in-phase signal after IQ demodulation. An additional circulator between cavity and JPA is used to further suppress leakage of the pump into the cavity (Fig. S1).

The ADwin samples the readout signal once, at a variable delay following a trigger from an arbitrary waveform generator (Tektronix AWG5014). This delay is adjusted in 100 ns steps to maximize readout contrast between states prepared with and without a π pulse on the $0 \leftrightarrow 1$ transition. For multiple feedback instances, the timing for each measurement is separately optimized. Once these timings are set, a routine (repeated every ~ 5 min), determines the optimum threshold for digitizing the readout signal. This voltage is then used to assign H or L to the measurement in any feedback cycle. For Fb_0 (Fb_1), the ADwin triggers another arbitrary waveform generator (Tektronix AWG520) to produce a π pulse when the outcome is L (H). Pulse timings and signal delays in the feedback cycle are illustrated in Fig. S2.

MEASUREMENT OF THE TRANSMON POPULATIONS

Whereas readout for feedback maximally discriminates between $|0\rangle$ and $|1\rangle$ from single measurements, the final readout, used to quantify the reset error $P_{\rm err}$, relies on the average of thousands of runs. Due to leakage from the qubit subspace, the output voltage $V_{\rm H}$ also depends on $P_{|2\rangle}$. The average measurement is related to the populations as [3]

$$\langle V_{\rm H} \rangle = \langle V_{|0\rangle} \rangle P_{|0\rangle} + \langle V_{|1\rangle} \rangle P_{|1\rangle} + \langle V_{|2\rangle} \rangle P_{|2\rangle}, \qquad \text{(S1)}$$

where the $\langle V_{|i\rangle} \rangle$ are the mean values of the distributions in Fig. S3. When $\langle V_{|1\rangle} \rangle \neq \langle V_{|2\rangle} \rangle$, as in Fig. S3(a), the measurement of $\langle V_{\rm H} \rangle$ is not sufficient to extract the three populations. Two possible strategies to obtain two additional equations are: 1) repeating the sequence, but swapping two of the populations prior to measurement in order to obtain different $\langle V'_{\rm H} \rangle$, $\langle V''_{\rm H} \rangle$; 2) analyzing the full time-domain response of the cavity [3]. We use a different approach, taking advantage of the degree of freedom in the relative phase between RF and pump, to extract $P_{|0\rangle}$ with the minimum number of measurements. As shown in Fig. S3(b), the phase between RF and pump is chosen to maximize overlap between the histograms for $|1\rangle$ and $|2\rangle$. Making $\langle V_{|1\rangle} \rangle = \langle V_{|2\rangle} \rangle$, and assuming $\sum_{i=0}^{2} P_{|i\rangle} = 1$ (as justified below), simplifies Eq. (S1) to

$$\langle V_{\rm H} \rangle = P_{|0\rangle} \langle V_{|0\rangle} \rangle + (1 - P_{|0\rangle}) \langle V_{|1\rangle} \rangle, \qquad \text{(S2)}$$

with $P_{|0\rangle}$ as the only unknown. To measure $P_{|1\rangle}$, we transfer this population to $|0\rangle$ with a π pulse, making $\langle V_{\rm H} \rangle = P_{|1\rangle} \langle V_{|0\rangle} \rangle + (1 - P_{|1\rangle}) \langle V_{|1\rangle} \rangle$. Similarly, to measure $P_{|2\rangle}$, we transfer this population to $|0\rangle$ with two consecutive π pulses, the first on $1 \leftrightarrow 2$ and the second on $0 \leftrightarrow 1$. The amplitudes $\langle V_{|0\rangle} \rangle$, $\langle V_{|1\rangle} \rangle$ and $\langle V_{|2\rangle} \rangle$ are calibrated by preparing the corresponding states. To do this, we initialize to $|0\rangle$ as best we can ($P_{|0\rangle} = 99.3\%$) by performing a measurement (used for postselection [1, 4]) and then applying the required π pulses. Values for $\langle V_{|0\rangle} \rangle$, $\langle V_{|1\rangle} \rangle$ and $\langle V_{|2\rangle} \rangle$ are estimated from $\sim 10\,000$ shots each. As shown in Fig. S3(b), $\langle V_{|1\rangle} \rangle$ and $\langle V_{|2\rangle} \rangle$ are matched to better than 1%. The sum of the extracted populations shown in Figs. S4(a-b) averages to 99.7% and 99.5% over the temporal evolutions starting from $|0\rangle$ and $|1\rangle$, respectively. For initial state $|2\rangle$ [Fig. S4(c)], leakage to the third excited state is relevant under 50 μs. In most cases, however, $P_{|2\rangle}$ is limited to a few percent, justifying the truncation of the model to the lowest three transmon levels.

MODEL FOR MEASUREMENT AND FEEDBACK

The readout errors ϵ_i^M, indicating the probability of a wrong result M for the pre-measurement state $|i\rangle$, are

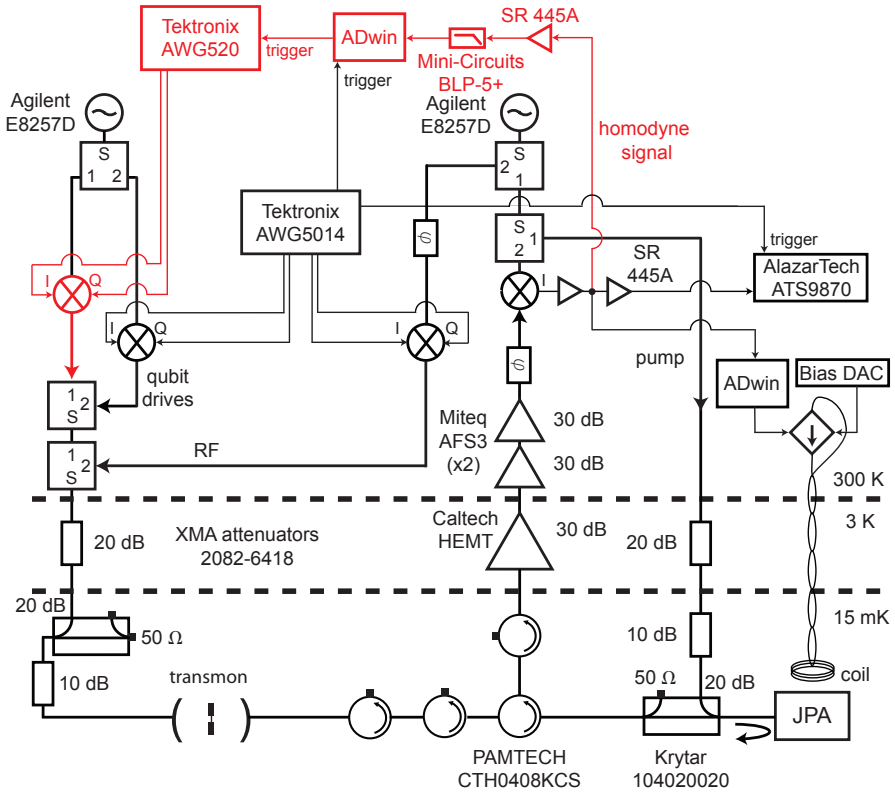

FIG. S1. Experimental setup. The key components of the feedback loop are highlighted in red. A Tektronix AWG5014 generates modulation envelopes for measurement and deterministic qubit pulses. The readout signal is amplified with a JPA at base temperature, a HEMT (Caltech CITCRYO 1-12) at 3 K, and two GaAs FET amplifiers (Miteq AFS3) at room temperature prior to demodulation to baseband (0 Hz intermediate frequency). The demodulated signal is amplified (Stanford Research Systems SR445A) and then split. One arm is sent through a low-pass filter (Mini-Circuits BLP-5+, DC-5 MHz passband) to an analog input of an ADwin Gold processor, and the other to the data acquisition card (AlazarTech). The ADwin compares the input voltage to the optimum threshold for discriminating between $|0\rangle$ and $|1\rangle$. Depending on the feedback logic programmed into it, the ADwin sends a trigger to the Tektronix AWG520 when the input is above/below the threshold. The triggered AWG520 produces the modulation envelope for a π pulse sent via the input line, closing the feedback loop.

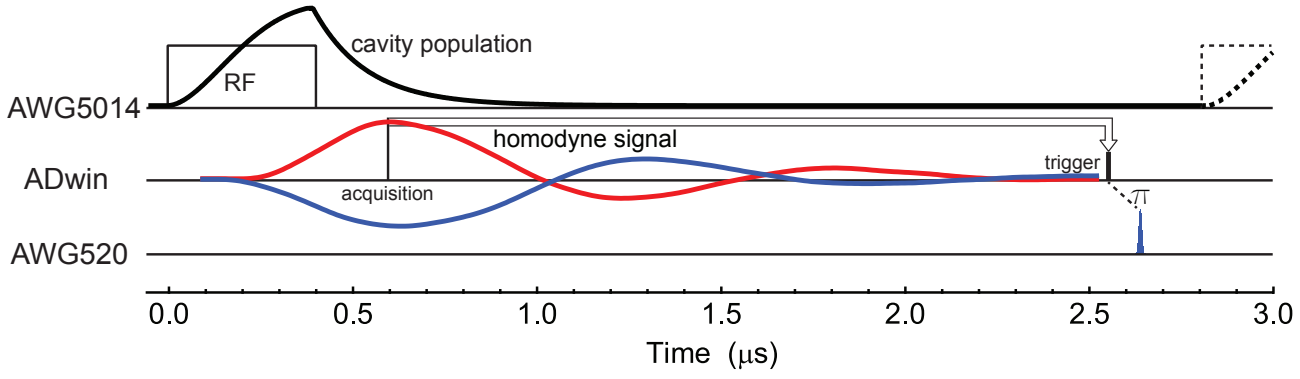

FIG. S2. Feedback loop latency. The AWG5014 modulates the RF tone to generate a 400 ns measurement pulse reaching the cavity at $t = 0$. The black curve is a sketch of the photon population inside the cavity, which reaches ~ 45 photons at $t = 0.4$ μs. We estimate a critical photon number $n_{\mathrm{crit}} \approx 150$ and inverse measurement rate $\Gamma_m^{-1} = 26$ ns at steady state [2]. The ADwin, triggered by the AWG5014, measures one channel of the output homodyne signal (red: qubit in $|0\rangle$, blue: $|1\rangle$), delayed by ~ 200 ns due to the low pass filter at its input side. After comparison of the measured voltage at $t = 0.6$ μs to the reference threshold, the AWG520 is conditionally triggered at $t = 2.54$ μs, resulting in a π pulse (Gaussian, $\sigma = 6$ ns, total length 24 ns), reaching the cavity at 2.62 μs. The ADwin acquisition (and the triggered π pulse) can be shifted by ± 0.1 μs to maximize readout fidelity. The measurement pulse for another feedback cycle or final readout reaches the cavity at $t = 2.8$ μs.

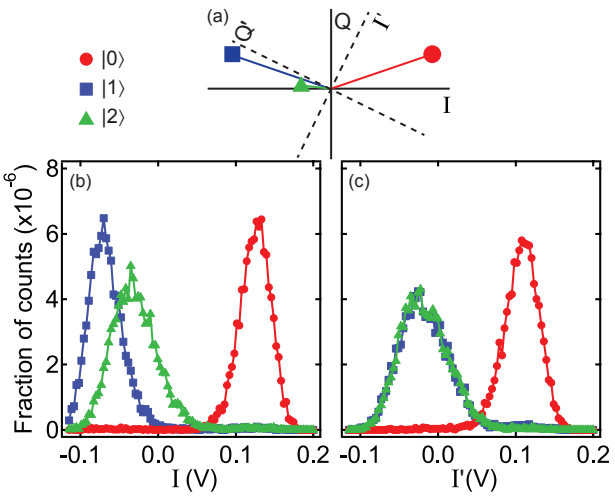

FIG. S3. The two readout configurations employed and the associated histograms for the three lowest-energy transmon levels. (a) Sketch in phase space of the coherent states corresponding to the first three transmon levels. The amplified quadrature, chosen by tuning the RF-pump phase difference, is either I or I′, giving the histograms in (b) and (c), respectively. (b) Optimum discrimination between $|0\rangle$ and $|1\rangle$. This choice is used in M_0 and in all the measurements which are part of a feedback sequence to maximize the single-shot readout fidelity. (b) Readout configuration making $|1\rangle$ and $|2\rangle$ indistinguishable, as their respective coherent states have equal projection onto the I′ axis. This configuration is used in the final measurement of every sequence to accurately extract the level populations, as described in the text.

extracted from the histograms in Figs. 2 and S3, using the result of the fit in Fig. S4 to account for the transmon equilibration between M_0 and M_1 [1]. The data yield the values $\epsilon_0^L = 0.3\%, \epsilon_1^H = 0.7\%$. By analogy with Ref. 1, we approximate the errors with $\epsilon_0^L \sim p_{00}^L$ and $\epsilon_1^H \sim p_{01}^H$ (notation defined in the main text). Combining these with the best-fit values for Fig. S3, $p_{00}^L + p_{10}^H = 0.4\%$ and $p_{11}^H + p_{01}^L = 1.4\%$, we estimate the remaining errors $p_{10}^H = 0.1\%$, $p_{01}^L = 1.0\%$.

COMPARISON OF FEEDBACK PROTOCOLS

A double round of feedback tends to equalize the initialization error for different input states (Figs. 3 and S5). Augmenting the protocol to three or more rounds only marginally affects the residual difference, which is mainly due to leakage into $|2\rangle$. When a π pulse is repeatedly applied on the $0 \leftrightarrow 1$ transition [Figs. 3(a) and S5(a)], this leakage depends on the fraction of time that $|1\rangle$ is populated, which increases at shorter $\tau_{\rm init}$. Therefore, $P_{|2\rangle}$ can exceed its steady-state value, leading to increased $P_{\rm err}$. We address this problem by inverting the populations in $|1\rangle$ and $|2\rangle$ after two iterations of Fb_0. This brings the population, previously in $|2\rangle$, within reach of a third instance of Fb_0. This strategy is only effective when $P_{|2\rangle} > P_{|1\rangle}$, explaining why a residual error of $\sim 3\%$ is observed at long $\tau_{\rm init}$, when, after two rounds of feedback, $P_{|1\rangle} \approx P_{|2\rangle}$. Similarly, the repeated initialization in $|0\rangle$ without periodically applying a π pulse [Figs. 4(b) and S5(b)] limits the escape probability to the excited states, as evident from the reduced error at short $\tau_{\rm init}$.

[1] D. Ristè, J. G. van Leeuwen, H.-S. Ku, K. W. Lehnert, and L. DiCarlo, arXiv:1204.2479, Phys. Rev. Lett., in press (2012).

[2] J. Gambetta, A. Blais, D. I. Schuster, A. Wallraff, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Phys. Rev. A **74**, 15 (2006).

[3] R. Bianchetti, S. Filipp, M. Baur, J. Fink, C. Lang, L. Steffen, M. Boissonneault, A. Blais, and A. Wallraff, Phys. Rev. Lett. **105**, 223601 (2010).

[4] J. E. Johnson, C. Macklin, D. H. Slichter, S. J. Weber, E. B. Weingarten, J. Clarke, and I. Siddiqi, arXiv:1202.5541, Phys. Rev. Lett., in press (2012).

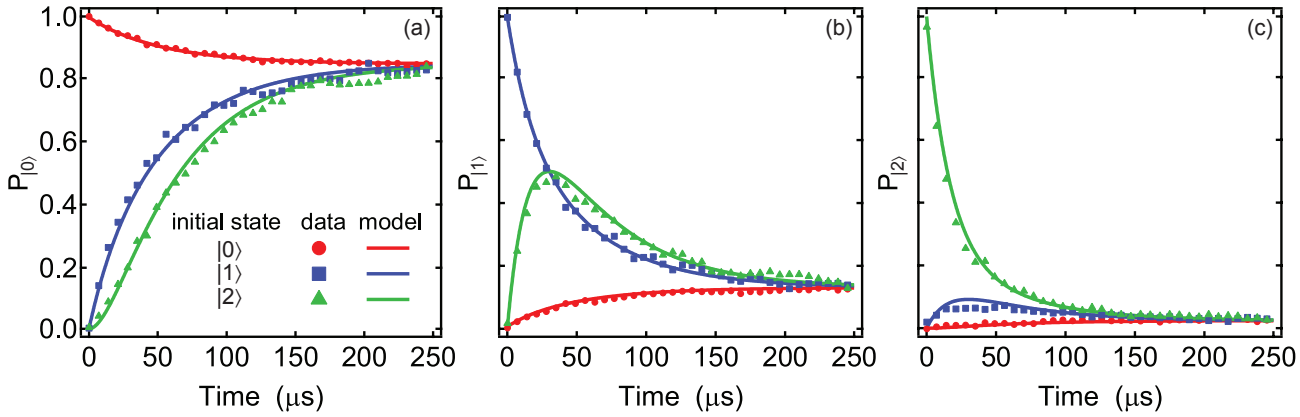

FIG. S4. Transmon population dynamics showing residual excitation. Populations (a) $P_{|0\rangle}$, (b) $P_{|1\rangle}$, and (c) $P_{|2\rangle}$ as a function of time following preparation in $|0\rangle$ (red), $|1\rangle$ (blue), and $|2\rangle$ (green). The combined fit of the master equation model to these data gives the transition rates and steady-state populations shown in Fig. 1.

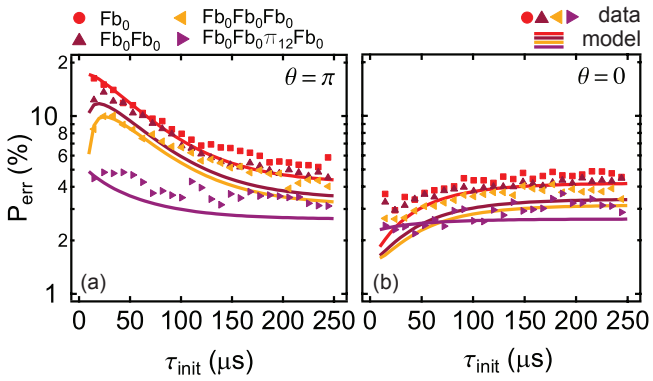

FIG. S5. Fast reset using multiple instances of Fb_0. The measurement sequence is illustrated in Fig. 4, with (a) $\theta = \pi$, and (b) $\theta = 0$. At any $\tau_{\rm init}$ $P_{\rm err}$ decreases by a factor ~ 3 with one instance of Fb_0 compared to passive initialization (Fig. 4). Two or three iterations add only a modest improvement. An additional decrease in error by a factor ~ 2 at short $\tau_{\rm init}$ is obtained with a π pulse on the $1 \leftrightarrow 2$ transition preceding the third instance (purple). This pulse corrects for the leakage of population to $|2\rangle$, occurring for short $\tau_{\rm init}$. The last feedback round brings this population back to $|0\rangle$. Solid curves are given by the model, using the parameters extracted from Figs. 1 and 2 and assuming perfect pulses.